

An improved approximation algorithm for the minimum cost subset k -connected subgraph problem

Bundit Laekhanukit^{*†}

January 21, 2013

Abstract

The minimum cost subset k -connected subgraph problem is a cornerstone problem in the area of network design with vertex connectivity requirements. In this problem, we are given a graph $G = (V, E)$ with costs on edges and a set of terminals T . The goal is to find a minimum cost subgraph such that every pair of terminals are connected by k openly (vertex) disjoint paths. In this paper, we present an approximation algorithm for the subset k -connected subgraph problem which improves on the previous best approximation guarantee of $O(k^2 \log k)$ by Nutov (FOCS 2009). Our approximation guarantee, $\alpha(|T|)$, depends upon the number of terminals:

$$\alpha(|T|) = \begin{cases} O(k \log^2 k) & \text{if } 2k \leq |T| < k^2 \\ O(k \log k) & \text{if } |T| \geq k^2 \end{cases}$$

So, when the number of terminals is *large enough*, the approximation guarantee improves significantly. Moreover, we show that, given an approximation algorithm for $|T| = k$, we can obtain almost the same approximation guarantee for any instances with $|T| > k$. This suggests that the hardest instances of the problem are when $|T| \approx k$.

1 Introduction

We present an improved approximation algorithm for the *minimum cost subset k -connected subgraph problem*. In this problem (subset k -connectivity, for short), we are given a graph $G = (V, E)$ with edge costs and a set of terminals $T \subseteq V$. The goal is to find a minimum cost subgraph such that each pair of terminals is connected by k openly (vertex) disjoint paths. This is a fundamental problem in network design which includes as special cases the minimum cost Steiner tree problem (the case $k = 1$) and the minimum cost k vertex-connected spanning subgraph problem (the case $T = V$). However, the subset k -connectivity problem is significantly harder than these two special cases. Specifically, an important result of Kortsarz, Krauthgamer and Lee [14] shows that the problem does not admit an approximation guarantee better than $2^{\log^{1-\epsilon} n}$ for any $\epsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}(n^{O(\text{polylog}(n))})$. In contrast, polylogarithmic approximation guarantees are known

^{*}School of Computer Science, McGill University. Email: blaekh@cs.mcgill.ca.

[†]This research was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) grant no. 288334, 429598 and by Harold H Helm fellowship.

[‡]Most of the work was done while the author was in Department of Combinatorics and Optimization, University of Waterloo, Canada.

for the minimum cost k -vertex connected spanning subgraph problem. The first such result was obtained by Fakcharoenphol and Laekhanukit [9] using the *Halo-set decomposition*, introduced by Kortsarz and Nutov [15]. Subsequently, Nutov [17] improved the approximation guarantee to $O\left(\log k \cdot \log \frac{n}{n-k}\right)$.

Since the hardness result of Kortsarz et al. [14] no non-trivial approximation algorithm was known for the general case of the subset- k -connectivity problem until the work of Chakraborty, Chuzhoy and Khanna [3]. They presented an $O(k^{O(k^2)} \cdot \log^4 |T|)$ -approximation algorithm for the rooted version of our problem, namely the *rooted subset k -connectivity problem*. There, given a root vertex r and a set of terminals T , the goal is to find a minimum cost subgraph that has k openly disjoint paths from the root vertex r to every terminal in T . Chakraborty et al. showed how to solve the subset k -connectivity problem by applying the rooted subset k -connectivity algorithm k times, thus obtaining an $O(k^{O(k^2)} \cdot \log^4 |T|)$ -approximation algorithm. Recently, in a series of developments [3, 7, 8, 21, 18], the approximation guarantees for the rooted subset k -connectivity problem has been steadily improved. This has culminated in an $O(k \log k)$ guarantee due to Nutov [21], thus implying an approximation guarantee of $O(k^2 \log k)$ for the subset k -connectivity problem.

There is a trivial way to obtain an approximation bound of $O(|T|^2)$. So, with the current progress on the rooted subset k -connectivity problem, the application of the rooted subroutine is only useful when the number of terminals is large enough, say $|T| \geq 2k$. The main contribution of this paper is to show that, in this case, only a polylogarithmic number of applications of the rooted subset k -connectivity algorithm are required to solve the subset k -connectivity problem. Given an approximation algorithm for the rooted subset k -connectivity problem, we show that only $O(\log^2 k)$ applications of the algorithm are required, and we can save a factor of $O(\log k)$ since some of these applications are applied to instances with lower costs. Moreover, as the number of terminal increases above k^2 , we are able to save an additional $O(\log k)$ factor. Thus, given an approximation algorithm for the rooted subset k -connectivity problem in [21] due to Nutov (and with careful analysis), we achieve an $\alpha(|T|)$ -approximation guarantee where

$$\alpha(|T|) = \begin{cases} O(k \log^2 k) & \text{if } 2k \leq |T| < k^2 \\ O(k \log k) & \text{if } |T| \geq k^2 \end{cases}$$

As we may combine our algorithm with the trivial $O(|T|^2)$ -approximation algorithm for the case $|T| < 2k$, we obtain an approximation guarantee of $O(k^2)$, which improves upon the previous best approximation guarantee of $O(k^2 \log k)$ for all cases. Moreover, for $|T| \geq 2k$, we obtain a significant improvement of a factor of k . Observe, however, that for the case $|T| \approx k$ the guarantee is still quadratic. At first, this may seem paradoxical since we may hope that the problem is easier when the number of terminals is small. Our results suggest that this is not the case. Indeed, it appears that the hardest instances of subset k -connectivity may have at most k terminals. Precisely, we show that, given an $\alpha(k)$ -approximation algorithm for the subset k -connectivity problem with $|T| = k$, there is an $(\alpha(k) + f(k))$ -approximation algorithm for any instance with $|T| > k$, where $f(k)$ is the best known approximation guarantee for the rooted subset k -connectivity problem. Furthermore, we give an approximation preserving reduction from the rooted subset k -connectivity problem to the subset k -connectivity problem, showing a strong connection between the two problems.

Related Work. Some very special cases of the subset k -connectivity problem are known to have constant factor approximation algorithms. For $k = 1$, the minimum cost Steiner tree problem, the

best known approximation guarantee is 1.39 due to Byrka, Grandoni, Rothvoß and Sanità [2]. For $k = 2$, a factor two approximation algorithm was given by Fleischer, Jain and Williamson [10]. The subset k -connectivity problem also has an $O(1)$ -approximation algorithm when edge costs satisfy the triangle inequality; see Cheriyan and Vetta [6]. The most general problem in this area is the vertex-connectivity survivable network design problem (VC-SNDP). In VC-SNDP, the connectivity requirement for each pair of vertices can be arbitrary. Recently, Chuzhoy and Khanna [8] showed that there is an $O(k^3 \log n)$ -approximation algorithm for VC-SNDP. The problems where requirements are edge and element connectivity (EC-SNDP and Element-SNDP) are also very well studied. Both problems admit 2-approximation algorithms via iterative rounding. For EC-SNDP, a 2-approximation algorithm was given by Jain [12]. For Element-SNDP a 2-approximation algorithm was given by Fleischer, Jain and Williamson [10]. The vertex-cost versions of these problems have also been studied in literature. Nutov [19] gave an approximation guarantee of $O(k \log |T|)$ for vertex-cost EC-SNDP using a technique, called spider decomposition. Later on, in [21], Nutov applied the spider decomposition technique to other vertex-cost problems, giving approximation guarantees of $O(k \log |T|)$ for Element-SNDP, $O(k^2 \log |T|)$ for the rooted subset k -connectivity problem, $O(k^3 \log |T|)$ for the subset k -connectivity problem and $O(k^4 \log^2 |T|)$ for VC-SNDP.

2 Preliminaries and Results

We begin with some formal definitions. Let $G = (V, E)$ denote the graph for an instance of the problem. For a set of edges F , the graph $G' = (V, E \cup F)$ is denoted by $G + F$; for a vertex v , the graph obtained from G by removing v is denoted by $G - v$. For any set of vertices $U \subseteq V$, let $\Gamma(U)$ denote the set of *neighbors* of U ; that is, $\Gamma(U) = \{v \in V - U : \exists (u, v) \in E, u \in U\}$. Define a set U^* to be $V - (U \cup \Gamma(U))$, which is the *vertex-complement* of U . For any pair of vertices $s, t \in V$, two s, t -paths are *openly disjoint* if they have no vertices except s and t in common. Let $T \subseteq V$ be a set of vertices called *terminals*. Without loss of generality, assume that no two terminals of T are adjacent in G . This assumption can be easily justified by subdividing every edge joining two terminals; that is, if there is an edge (s, t) joining two terminals, then we replace (s, t) by two new edges (s, u) and (u, t) and set costs of the new edges so that $c(s, t) = c(s, u) + c(u, t)$, where $c(\cdot)$ is a cost function. The set of terminals T is *k -connected in G* if the graph G has k openly disjoint s, t -paths between every pair of terminals $s, t \in T$. Thus, by Menger Theorem, the removal of any set of vertices with size at most $k - 1$ leaves all the remaining terminals in the same component of the remaining graph. By the *subset connectivity* of G on T , we mean the maximum integer ℓ such that T is ℓ -connected in G . A *deficient set* is a subset of vertices $U \subseteq V$ such that both U and U^* contain terminals of T and $|\Gamma(U)| < k$. Observe that the vertex-complement U^* is also a deficient set. Similarly, given a designated *root* vertex r , the graph is *k -connected from r to T* if G has k openly disjoint r, t -paths for every terminal $t \in T$ (r may or may not be in T). By the *rooted connectivity* of G from r to T , we mean the maximum integer ℓ such that G is ℓ -connected from r to T .

In the *subset k -connectivity problem*, we are given a graph $G = (V, E)$ with a cost $c(e)$ on each edge $e \in E$, a set of terminals $T \subseteq V$, and an integer $k \geq 0$. The goal is to find a set of edges $\hat{E} \subseteq E$ of minimum cost such that T is k -connected in the subgraph $\hat{G} = (V, \hat{E})$. In the *rooted subset k -connectivity problem*, our goal is to find a set of edges $\hat{E} \subseteq E$ of minimum cost such that the subgraph $\hat{G} = (V, \hat{E})$ is k -connected from r to T , for a given root r .

Nutov [21] recently gave an $O(k \log k)$ -approximation algorithm for the rooted subset k -connectivity

problem. The approximation guarantee improves by a logarithmic factor for the problem of increasing the rooted connectivity of a graph by one.

Theorem 2.1 (Nutov 2009 [21]). *There is an $O(k \log k)$ -approximation algorithm for the rooted subset k -connectivity problem. Moreover, consider the restricted version of the problem where the goal is to increase the rooted connectivity from ℓ to $\ell + 1$. Then the approximation guarantee (with respect to a standard LP) is $O(\ell)$.*

Our focus is upon the subset k -connectivity problem. The followings are our main results:

Theorem 2.2. *For any set T of terminals, there is an $\alpha(|T|)$ -approximation algorithm for the subset k -connectivity problem where*

$$\alpha(|T|) = \begin{cases} O(|T|^2) & \text{if } |T| < 2k \quad (\text{folklore}) \\ O(k \log^2 k) & \text{if } 2k \leq |T| < k^2 \\ O(k \log k) & \text{if } |T| \geq k^2 \end{cases}$$

In particular, there is an $O(k^2)$ -approximation algorithm for the general case of the subset k -connectivity problem, and there is an $O(k \log k)$ -approximation algorithm when $|T| \geq k^2$.

Proposition 2.3. *Consider the subset k -connectivity problem. Suppose there is an $\alpha(k)$ -approximation algorithm for instances with $|T| = k$. Then there is an $(\alpha(k) + f(k))$ -approximation algorithm for any instance with $|T| > k$, where $f(k)$ is the best known approximation guarantee for the rooted subset k -connectivity problem.*

Theorem 2.4. *There is an approximation preserving reduction such that, given an instance of the rooted subset k -connectivity problem consisting of a graph G , a root vertex r and a set of terminals T , outputs an instance of the subset k -connectivity problem consisting of a graph G' and a set of terminals $T \cup \{r\}$.*

The hardness result in Theorem 2.4 together with the hardness of the rooted subset k -connectivity problem by Cheriyan, Laekhanukit, Naves and Vetta [4] implies the hardness of $\Omega(k^\epsilon)$, for the subset k -connectivity problem, where $\epsilon > 0$ is some fixed constant.

Some results and proofs similar to the ones in this paper have appeared in previous literature; see [5, 15, 16]. In particular, Lemma 3.10 and Lemma 3.13 appeared in [15] and [16], respectively. The proofs of Proposition 2.3 is identical to that of the case $T = V$, which was given in [13] and also in [1]. Our key new contributions are Lemmas 3.5, 3.9 and 3.8, which allow us to extend the result in [15] to the subset k -connectivity problem.

We remark that, at the time this paper is written, the approximation guarantee of the subset k -connectivity problem was improved by Nutov [20] to $O(k \log k)$ for all $k \leq |T| - o(|T|)$.

Organization: In Section 3, we present an approximation algorithm for the subset k -connectivity problem, which is the main result in this paper. In Section 3.7, we give a discussion that our algorithm and analysis can be extended to the case $k < |T| < 2k$. To keep the presentation simple, Section 3.7 is presented separately from the main result. In Section 4, we discuss the hardness of the subset k -connectivity problem. To be precise, we show that the hardest instance of the subset k -connectivity problem might be when $|T| \approx k$, and we give an approximation preserving reduction from the rooted subset k -connectivity problem to the subset k -connectivity problem.

3 An approximation algorithm

Our main result in Theorem 2.2 breaks up into three cases where there are a small number, a moderate number and a large number of terminals, respectively. Indeed, the first case is a folklore. When there are a small number of terminals ($|T| < 2k$), we apply the following trivial $O(|T|^2)$ -approximation algorithm. We find k openly disjoint paths of minimum cost between every pair of terminals by applying a minimum cost flow algorithm. Let opt denote the cost of the optimal solution to the subset k -connectivity problem. Since any feasible solution to the subset k -connectivity problem has k openly disjoint paths between every pair of terminals, the cost incurred by finding a minimum cost collection of k openly disjoint paths between any pair of terminals is at most opt . Since we have at most $|T|^2$ pairs, this incurs a total cost of $O(|T|^2 \cdot opt)$.

The remaining two cases are similar. Things are slightly easier, though, when there are large number of terminals ($|T| \geq k^2$), leading to a slightly better guarantee than when there are a moderate number of terminals ($2k \leq |T|$). We devote most of this section to presenting an approximation algorithm for the moderate case. (In Section 3.6, we show the improvement for the case of a large number of terminals.)

Our algorithm works by repeatedly increasing the subset connectivity of a graph by one. We start with a graph that has no edges. Then we apply k outer iterations. Each outer iteration increases the subset connectivity (of the current graph) by one by adding a set of edges of approximately minimum cost. The analysis of the outer iterations applies linear programming (LP) scaling and incurs a factor of $O(\log k)$ in the approximation guarantee for the k outer iterations. The analysis based on LP-scaling can be seen in [22, 6, 15, 9] and also in [11, 16].

The following is a standard LP-relaxation for the subset k -connectivity problem.

$$\begin{aligned} \min \quad & \sum_{e \in E} c_e x_e \\ \text{s.t.} \quad & \sum_{e \in \delta(U, W)} x_e \geq k - |V - (U \cup W)| \quad \forall (U, W) \in \mathcal{S} \\ & 0 \leq x_e \leq 1 \quad \forall e \in E \end{aligned}$$

where $\delta(U, W) = \{(u, w) \in E : u \in U, w \in W\}$ is a set of edges with one endpoint in U and the other endpoint in W , and $\mathcal{S} = \{(U, W) \in V \times V : U \cap W = \emptyset, U \cap T \neq \emptyset, W \cap T \neq \emptyset\}$.

Lemma 3.1. *Suppose there is a $\beta(\ell)$ -approximation algorithm for the problem of increasing the subset connectivity of a graph from ℓ to $\ell + 1$ with respect to a standard LP, where $\beta(\ell)$ is a non-decreasing function. Then there is an $O(\beta(k) \log k)$ -approximation algorithm for the subset k -connectivity problem.*

We are left with the key problem of increasing the subset connectivity (of the current graph) by one by adding a set of edges of approximately minimum cost. Throughout this section, we assume that the set of terminals T is ℓ -connected in the current graph, and $|T| \geq 2k \geq 2\ell$. Also, we assume that no two terminals are adjacent in the input graph $G = (V, E)$.

Assumption: The set of terminals T is ℓ -connected in the current graph $\widehat{G} = (V, \widehat{E})$. Moreover, no two terminals are adjacent in the input graph G .

Our algorithm solves the problem of increasing the subset connectivity of a graph by one by applying a number of so-called *inner* iterations. To describe our algorithm, we need some definitions

and subroutines. Thus, we defer the description of our algorithm to Section 3.5. In Section 3.1, we give important definitions and structures of subset ℓ -connected graphs called “cores” and “halo-families”. Our algorithm requires two subroutines. The first one is the subroutine that employs the rooted subset $(\ell + 1)$ -connectivity algorithm to cover halo-families. This subroutine is given in Section 3.2. The second one is the subroutine for decreasing the number of cores to $O(\ell)$, which is given in Section 3.3. Then we introduce a notion of “thickness” in Section 3.4. This notion guides us how to use the rooted subset $(\ell + 1)$ -connectivity algorithm efficiently. Finally, in Section 3.5, we present an $O(k \log^2 k)$ -approximation algorithm for the case $|T| \geq 2k$. By slightly modifying the algorithm and analysis, we show in Section 3.6 that our algorithm achieves a better approximation guarantee of $O(k \log k)$ when $|T| \geq k^2$.

3.1 Subset ℓ -connected graphs: deficient sets, cores, halo-families and halo-sets

In this section, we discuss some key properties of deficient sets that will be exploited by our approximation algorithm.

Assume that the set of terminals T is ℓ -connected in the graph $G = (V, E)$. Then G has $|\Gamma(U)| \geq \ell$ for all $U \subseteq V$ such that $U \cap T \neq \emptyset$ and $U^* \cap T \neq \emptyset$. Moreover, by Menger Theorem, G is subset $(\ell + 1)$ -connected if and only if G has no deficient set.

A key property of vertex neighborhoods is that the function $|\Gamma(\cdot)|$ on subsets of V is submodular. In other words, for any subsets of vertices $U, W \subseteq V$,

$$|\Gamma(U \cup W)| + |\Gamma(U \cap W)| \leq |\Gamma(U)| + |\Gamma(W)|.$$

We call a deficient set $U \subseteq V$ *small*¹ if $|U \cap T| \leq |U^* \cap T|$.

Proposition 3.2. *For any small deficient set U , $|U \cap T| \leq |T|/2$ and $|U^* \cap T| \geq (|T| - \ell)/2$.*

Proof. The first inequality follows from the definition of small deficient sets. Consider the second inequality. We have

$$|U^* \cap T| \geq \frac{|U \cap T| + |U^* \cap T|}{2} = \frac{|T| - |\Gamma(U) \cap T|}{2} \geq \frac{|T| - \ell}{2}$$

□

Lemma 3.3 (Uncrossing Lemma). *Consider any two distinct deficient sets $U, W \subseteq V$. If $U \cap W \cap T \neq \emptyset$ and $U^* \cap W^* \cap T \neq \emptyset$, then both $U \cap W$ and $U \cup W$ are deficient sets. Moreover, if U or W is a small deficient set, then $U \cap W$ is a small deficient set.*

Proof. Suppose $U \cap W \cap T \neq \emptyset$ and $U^* \cap W^* \cap T \neq \emptyset$. Note that

$$\begin{aligned} U^* \cap W^* \cap T &= (V - (U \cup \Gamma(U))) \cap (V - (W \cup \Gamma(W))) \cap T \\ &= (V - (U \cup W \cup \Gamma(U) \cup \Gamma(W))) \cap T \\ &= (V - ((U \cup W) \cup \Gamma(U \cup W))) \cap T \\ &= (U \cup W)^* \cap T. \end{aligned}$$

¹ There is another way to define a *small* deficient set. For example, in [20], Nutov defined a small deficient set as a deficient set U such that $|U \cap T| \leq \frac{|T| - \ell}{2}$.

Moreover, $(U \cup W)^* \subseteq (U \cap W)^*$. This means that

$$(U \cap W) \cap T \neq \emptyset, \quad (U \cup W) \cap T \neq \emptyset, \quad (U \cup W)^* \cap T \neq \emptyset \quad \text{and} \quad (U \cap W)^* \cap T \neq \emptyset.$$

Hence, by Menger Theorem, we have $|\Gamma(U \cup W)| \geq \ell$ and $|\Gamma(U \cap W)| \geq \ell$. Moreover, since U, W are deficient sets, we have $|\Gamma(U)| = |\Gamma(W)| = \ell$. It then follows by the submodularity of $|\Gamma(\cdot)|$ that

$$2\ell \leq |\Gamma(U \cup W)| + |\Gamma(U \cap W)| \leq |\Gamma(U)| + |\Gamma(W)| = 2\ell.$$

Thus, $|\Gamma(U \cap W)| = |\Gamma(U \cup W)| = \ell$. This implies that both $U \cup W$ and $U \cap W$ are deficient sets. Moreover, suppose U or W is a small deficient set. Without loss of generality, assume that U is a small deficient set. Then $U \cap W$ is a small deficient set because $U \cap W \subseteq U$. Thus,

$$|U \cap W \cap T| \leq |U \cap T| \leq |U^* \cap T| \leq |(U \cap W)^* \cap T|$$

□

By a *core*, we mean a small deficient set C that is inclusionwise minimal. In other words, C is a core if it is a small deficient set that does not contain another such set. It can be seen that any small deficient set U contains at least one core.

The *halo-family* of a core C , denoted by $\text{Halo}(C)$, is the set of all small deficient sets that contain C and contain no other cores; that is,

$$\text{Halo}(C) = \{U : U \text{ is a small deficient set, } C \subseteq U, \text{ and there is no core } D \neq C \text{ such that } D \subseteq U\}.$$

The *halo-set* of a core C , denoted by $H(C)$, is the union of all the sets in $\text{Halo}(C)$; that is,

$$H(C) = \bigcup \{U : U \in \text{Halo}(C)\}$$

An example of cores, halo-families and halo-sets is illustrated in Figure 1.

Remark: We remark that cores and halo-sets of subset ℓ -connected graphs can be computed in polynomial time. In fact, algorithms for computing cores and halo-sets of the k -vertex connected spanning subgraph problem also apply to the subset k -connectivity problem. See [15, 9, 16].

Some important properties of cores and halo-families that we will require are stated below.

Lemma 3.4 (Disjointness Lemma). *Consider any two distinct cores C and D . For any deficient sets $U \in \text{Halo}(C)$ and $W \in \text{Halo}(D)$, either $U \cap W \cap T = \emptyset$ or $U^* \cap W^* \cap T = \emptyset$.*

Proof. Suppose to the contrary that $U \cap W \cap T \neq \emptyset$ and $U^* \cap W^* \cap T \neq \emptyset$. Then, by Lemma 3.3, $U \cap W$ is a small deficient set. Thus, $U \cap W$ contains a core. This core is either C or D or another core distinct from C and D . In each case, we have a contradiction. □

The next result gives an upper bound on the number of halo-sets that contain a chosen terminal, which is a key for the design of our algorithm.

Lemma 3.5 (Upper bound). *For any terminal $t \in T$, the number of cores C such that $t \in H(C)$ is at most $\frac{2(|T|-1)}{|T|-\ell}$.*




Figure 1: The figure shows an example of cores, halo-families and halo-sets for $\ell = 1$. In the figure, we are given a graph $\widehat{G} = (V, \widehat{E})$. The graph \widehat{G} is a tree constructed by adding a vertex r as a root. Then we add terminals t_1, t_2, t_3, t_4 and add paths (t_1, v_1, r) , (t_2, v_1, r) , (t_3, v_3, v_2, r) and (t_4, v_4, r) connecting terminals to the root vertex. Clearly, $T = \{t_1, t_2, t_3, t_4\}$ is ℓ -connected in \widehat{G} . Every leaf vertex of the tree \widehat{G} forms a core. For example, $C = \{t_3\}$ is a core. The halo-family of C is $\text{Halo}(C) = \{\{t_3\}, \{t_3, v_3\}, \{t_3, v_2, v_3\}\}$, and the halo-set of C is $H(C) = \{t_3, v_2, v_3\}$. Observe that the set $\{t_1, t_2\}$ is a small deficient set, but it does not belong to any halo-family.

Proof. Let C_1, C_2, \dots, C_q be distinct cores such that $t \in H(C_i)$ for all $i = 1, 2, \dots, q$. For each $i = 1, 2, \dots, q$, since t is in the halo-set $H(C_i)$, there must exist a deficient set U_i in the halo-family $\text{Halo}(C_i)$ that contains t . It then follows that $t \in \bigcap_{i=1}^q U_i$. By the Disjointness Lemma (Lemma 3.4), for $i \neq j$, $U_i \cap U_j \cap T \neq \emptyset$ only if $U_i^* \cap U_j^* \cap T = \emptyset$. This is because U_i and U_j are deficient sets of different halo-families. For $i = 1, 2, \dots, q$, observe that U_i is small. Hence, by Proposition 3.2, we have $|U_i^* \cap T| \geq (|T| - \ell)/2$. Thus, the upper bound on the number halo-sets that contain t is

$$\frac{|T| - 1}{(|T| - \ell)/2} = \frac{2(|T| - 1)}{|T| - \ell}$$

□

3.2 Covering halo-families via rooted subset $(\ell + 1)$ -connectivity

We say that an edge $e = (u, v)$ covers a deficient set U if e connects U and U^* ; that is, $u \in U$ and $v \in U^*$. Clearly, e covers U if e covers U^* . Observe that if e covers U , then after adding the edge e to the current graph, U is no longer a deficient set. Now, consider any core C . We say that a set of edges F covers the halo-family of C if each deficient set U in $\text{Halo}(C)$ is covered by some edge of F . For a terminal $r \in T$, we say that the terminal r hits the halo-family $\text{Halo}(C)$ if r is in C or r is in the vertex-complement of the halo-set of C ; that is, r hits $\text{Halo}(C)$ if $r \in C$ or $r \in H(C)^*$. For a set of terminals $S \subseteq T$, we say that S hits a halo-family $\text{Halo}(C)$ if there is a terminal $r \in S$ that hits $\text{Halo}(C)$. The following lemma shows that if r hits the halo-family $\text{Halo}(C)$, then we can find a set of edges F that covers $\text{Halo}(C)$ by applying the rooted subset $(\ell + 1)$ -connectivity algorithm with r as the root.

Lemma 3.6. *Consider a set of edges F whose addition to \widehat{G} makes the resulting graph $\widehat{G} + F$ $(\ell + 1)$ -connected from a terminal r to T . Let C be any core. If $r \in C$ or $r \in H(C)^*$, then F covers all deficient sets in the halo-family of C .*

Proof. Consider the graph $\widehat{G} + F$. By the construction, $\widehat{G} + F$ has $(\ell + 1)$ openly disjoint paths from r to every terminal of T . This means that F covers all deficient sets of \widehat{G} that contains r . If $r \in C$, then $r \in U$ for all deficient sets $U \in \text{Halo}(C)$. So, F covers $\text{Halo}(C)$. Similarly, If $r \in H(C)^*$, then $r \in U^*$ for all deficient sets $U \in \text{Halo}(C)$. So, again, F covers $\text{Halo}(C)$ because an edge $e \in F$ covers U^* if and only if e covers U , and the lemma follows. \square

3.3 Preprocessing to decrease the number of cores

In this section, we describe the preprocessing algorithm that decreases the number of cores to $O\left(\frac{\ell|T|}{|T|-\ell}\right)$. We apply the following *root padding algorithm* in the preprocessing step.

The root padding algorithm: The algorithm takes as an input a graph $G = (V, E)$ with the given edge costs, a subset of terminals $R \subseteq T$, and a connectivity parameter $\rho \leq |R|$. We construct a padded graph by adding a new vertex \widehat{r} and new edges of zero cost from \widehat{r} to each terminal of R . Then we apply the rooted subset ρ -connectivity algorithm to the padded graph with the set of terminals T and the root \widehat{r} . We denote a solution subgraph (of the padded graph) by $\widehat{G} = (V \cup \{\widehat{r}\}, F \cup \{(\widehat{r}, t) : t \in R\})$, where $F \subseteq E$. Then the algorithm outputs the subgraph (of the original graph) $\widehat{G} - \widehat{r} = (V, F)$. The following result shows that, in the resulting graph \widehat{G} , every deficient set contains at least one terminal of R .

Lemma 3.7 (root padding). *Suppose we apply the root padding algorithm as above, and it finds a subgraph $\widehat{G} - \widehat{r} = (V, F)$. Then every deficient set of $\widehat{G} - \widehat{r}$ (with respect to ρ -connectivity of T in $\widehat{G} - \widehat{r}$) contains at least one terminal of R .*

Proof. Observe that \widehat{G} has ρ openly disjoint \widehat{r}, t -paths, for all $t \in T$. Suppose $U \subseteq V$ is a deficient set of $\widehat{G} - \widehat{r}$ that contains none of terminals of R . Thus, U contains another terminal $t \in T - R$ and $|\Gamma(U)| < \rho$. Then $(\widehat{G} - \widehat{r}) - \Gamma(U)$ has no path between t and a terminal of R . This also holds for $\widehat{G} - \Gamma(U)$ because adding \widehat{r} and the edges from \widehat{r} to every terminals of R cannot give a path between t and a terminal of R . This is a contradiction since \widehat{G} should have ρ openly disjoint \widehat{r}, t -paths. \square

Remark: Consider an instance Π_{root} of the rooted subset k -connectivity problem that is obtained from an instance Π_{subset} of the subset k -connectivity problem by either (1) picking one of the terminals as the root or (2) applying the root padding algorithm using any k terminals of Π_{subset} . Then the cost of the optimal solution to Π_{root} is at most the cost of the optimal solution to Π_{subset} because any feasible solution to Π_{subset} gives a feasible solution to Π_{root} (but not vice-versa).

Next, recall that $|T| \geq 2\ell$. We apply the root padding algorithm in Lemma 3.7 to any subset R of $(\ell + 1)$ terminals with $\rho = (\ell + 1)$. By Theorem 2.1, this incurs a cost of at most $O((k \log k) \cdot \text{opt})$. Moreover, the algorithm adds a set of edges to the current graph such that every deficient set of the resulting graph contains at least one terminal of R . Thus, each core of the resulting graph contains at least one terminal of R . By Lemma 3.5, each terminal is in $O\left(\frac{|T|}{|T|-\ell}\right) = O(1)$ halo-sets. Hence, the number of cores in the resulting graph is at most $O(\ell)$. This gives the next result.

Lemma 3.8. *Given a subset ℓ -connected graph, where $|T| \geq 2\ell$, there is an $f(k)$ -approximation algorithm that decreases the number of cores to $O(\ell)$, where $f(k)$ is the best known approximation guarantee for the rooted subset k -connectivity problem.*

3.4 Thickness of terminals

Consider a graph \widehat{G} such that T is ℓ -connected in \widehat{G} . We define the *thickness* of a terminal $t \in T$ to be the number of halo-families $\text{Halo}(C)$ such that $t \in \Gamma(H(C))$. Thus, the thickness of a terminal t is $|\{\text{Halo}(C) : C \text{ is a core, } t \in \Gamma(H(C))\}|$.

The following lemmas show the existence of a terminal with low thickness.

Lemma 3.9. *For every core C , $|\Gamma(H(C))| \leq \ell$.*

Proof. We use induction on the number of deficient sets in $\text{Halo}(C)$. For the induction basis, $\text{Halo}(C)$ has one deficient set U , and $H(C) = U$. Then $\Gamma(H(C)) = \Gamma(U)$ has size ℓ since the graph is subset ℓ -connected.

Suppose that \mathcal{U} is the union of j deficient sets that each contains the core C , and suppose that $|\Gamma(\mathcal{U})| \leq \ell$. Consider another deficient set W that contains C . Our goal is to show that $\Gamma(\mathcal{U} \cup W)$ has size at most ℓ . If $W \subseteq \mathcal{U}$, then we are done. Otherwise, we apply the submodularity of $|\Gamma(\cdot)|$. Observe that $\mathcal{U} \cap W$ contains a terminal since $\mathcal{U} \cap W \supseteq C$, and $(\mathcal{U} \cap W)^*$ contains a terminal since $(\mathcal{U} \cap W)^* \supseteq W^*$ and W^* contains a terminal. Hence, $\Gamma(\mathcal{U} \cap W)$ has size at least ℓ . Thus, we have

$$2\ell \geq |\Gamma(\mathcal{U})| + |\Gamma(W)| \geq |\Gamma(\mathcal{U} \cup W)| + |\Gamma(\mathcal{U} \cap W)| \geq |\Gamma(\mathcal{U} \cup W)| + \ell.$$

This implies that $\Gamma(\mathcal{U} \cup W)$ has size at most ℓ , and the lemma follows. \square

The following lemma shows the existence of a terminal with low thickness.

Lemma 3.10. *Consider a subset ℓ -connected graph. Let q denote the number of halo-families. Then there exists a terminal $t \in T$ with thickness at most $\frac{\ell q}{|T|}$.*

Proof. Consider the following bipartite incidence graph B of terminals and halo-families: B has a vertex for each terminal and each halo-family, and it has an edge between a terminal t and a halo-family $\text{Halo}(C)$ if and only if $t \in \Gamma(H(C))$. The previous lemma shows that each halo-family is adjacent to at most ℓ terminals in B . Hence, B has at most ℓq edges. Therefore, B has a terminal that is adjacent to at most $\frac{\ell q}{|T|}$ cores; that is, there exists a terminal with the required thickness. \square

3.5 An $O(k \log^2 k)$ -approximation algorithm for $|T| \geq 2k$

In this section, we describe our approximation algorithm for the case of a moderate number of terminals. Recall that we solve the problem by iteratively increasing the subset connectivity of a graph by one. Initially, we apply the algorithm in Section 3.3 to decrease the number of core to $O(\ell)$. Then we apply inner iterations until all the deficient sets are covered. At the beginning of each inner iteration, we compute the cores and the halo-sets. Then we apply a *covering-procedure* to find a set of edges that covers all the computed halo-families. This completes one inner iteration. Note that an inner iteration may not cover all of the deficient sets because deficient sets that contain two or more of the initial cores (those computed at the start of the inner iteration) may not be covered. So, we have to repeatedly apply inner iterations until no core is present. See Algorithm 1.

We now describe the covering-procedure. The procedure first finds a set of terminals $S \subseteq T$ that hits all the computed halo-families. Then it applies the rooted subset $(\ell + 1)$ -connectivity algorithm (Theorem 2.1) from each terminal of S . Let F be the union of all edges found by the rooted subset $(\ell + 1)$ -connectivity algorithm. Then, by Lemma 3.6, F covers all the halo-families.

The key idea of our algorithm is to pick a terminal \widehat{r} with a minimum thickness. Observe that a halo-family $\text{Halo}(C)$ is not hit by \widehat{r} only if

Algorithm 1 An approximation algorithm for moderate and large number of terminals

```

for  $\ell = 0, 1, \dots, k - 1$  do                                     ▷ (outer iterations)
    (* Increase the subset connectivity of a graph by one. *)
    Decrease the number of cores to  $O(\ell)$ .
    while the number of cores is greater than 0 do                 ▷ (inner iterations)
        Compute cores and halo-sets.
        Apply a covering-procedure to cover all the halo-families
    end while
end for

```

- (1) its halo-set $H(C)$ has \hat{r} as a neighbor (that is, $\hat{r} \in \Gamma(H(C))$) or
- (2) its halo-set $H(C)$ contains \hat{r} , but its core C does not contain \hat{r} .

The number of halo-families $\text{Halo}(C)$ such that $\hat{r} \in \Gamma(H(C))$ may be large, but the number of halo-families whose halo-sets contain \hat{r} is $O(1)$, assuming that $|T| \geq 2\ell$. Hence, we only hit halo-families of the second case by picking one terminal from each core C whose halo-set contains \hat{r} . Thus, the number of terminals picked is $O(1)$. We call this a *micro* iteration. Then the remaining halo-families are the halo-families whose halo-sets have \hat{r} as a neighbor. We repeatedly apply micro iterations until we hit all of the halo-families computed at the start of the inner iteration.

To be precise, initially let $S = \emptyset$. In each micro iteration, we add to S a terminal \hat{r} of minimum thickness (with respect to halo-families that are not hit by S). Then, for each core C such that $\hat{r} \in H(C) - C$, we add to S any terminal in $C \cap T$. We repeatedly apply micro iterations until S hits all the halo-families. At the termination, we apply the rooted subset $(\ell + 1)$ -connectivity algorithm (Theorem 2.1) from each terminal of S , and we return all the set of edges found by the algorithm as an output. The covering-procedure is presented in Figure 2.

Algorithm 2 Covering-procedure

```

 $S \leftarrow \emptyset$ .
while some halo-family is not hit by  $S$  do                         ▷ (micro iteration)
    Add to  $S$  a terminal  $\hat{r}$  with a minimum thickness.
    for each halo-family  $\text{Halo}(C)$  (not hit by  $S$ ) such that  $\hat{r} \in H(C) - C$  do
        Add to  $S$  any terminal  $r \in C$ .
    end for
end while
for each terminals  $r$  in  $S$  do
    Apply the rooted subset  $(\ell + 1)$ -connectivity algorithm from  $r$ .
end for

```

3.5.1 Analysis

The feasibility of a solution directly follows from the condition of the inner iteration; that is, the inner iteration terminates when a current graph has no core. So, at the termination of the inner iteration, the resulting graph has no deficient set. Thus, the subset connectivity of the graph becomes $\ell + 1$. Applying the outer iteration k times, the final graph is then subset k -connected.

It remains to analyze the cost of the solution subgraph. First, we analyze the number of times that the covering-procedure applies the rooted subset $(\ell + 1)$ -connectivity algorithm. Then we analyze the total cost incurred by all inner iterations, which is the cost for increasing the subset connectivity of a graph by one. Finally, we apply Theorem 3.1 to analyze the final approximation guarantee.

Consider any micro iteration of the covering-procedure. By Lemma 3.5, \hat{r} is contained in at most $O(1)$ halo-sets, assuming that $|T| \geq 2\ell$. Hence, we have to apply the rooted subset $(\ell + 1)$ -connectivity algorithm $O(1)$ times.

We now analyze the number of micro iterations needed to hit all of the halo-families. Let h_i denote the number of halo-families that are not hit by S at the beginning of the i -th micro iteration. Recall that the number of cores after the preprocessing step is $O(\ell)$. Thus, $h_1 = O(\ell)$. We claim that, at the i -th iteration, the number of halo-families that are not hit by S is at most $h_1/2^{i-1}$.

Lemma 3.11. *Consider the i -th micro iteration. The number of halo-families that are not hit by terminals of S at the start of the iteration is $h_1/2^{i-1}$.*

Proof. We proceed by induction on i . It is trivial for $i = 1$. Suppose that the assertion is true for the $(i - 1)$ -th micro iteration for some $i > 1$. Consider the $(i - 1)$ -th micro iteration. Since we choose a vertex \hat{r} with a minimum thickness, by Lemma 3.10, the thickness of \hat{r} is at most $\frac{h_{i-1}\ell}{|T|}$. Note that $\ell/|T| \leq 1/2$ since $|T| \geq 2\ell$. This means that \hat{r} is a neighbor of at most $h_{i-1}/2$ halo-sets. At the end of the micro iteration, halo-families that are not hit by terminals of S are halo-families whose halo-sets have \hat{r} as a neighbor. Thus, the number of remaining halo-families is at most $h_{i-1}/2$. Hence, we have

$$h_i \leq h_{i-1} \left(\frac{1}{2}\right) \leq \left(h_1 \left(\frac{1}{2}\right)^{i-2}\right) \left(\frac{1}{2}\right) = h_1 \left(\frac{1}{2}\right)^{i-1}$$

□

Lemma 3.11 implies that the maximum number of micro iterations (within the covering-procedure) is $O(\log h_1) = O(\log \ell)$. So, in each inner iteration, we have to call the rooted subset $(\ell + 1)$ -connectivity algorithm $O(\log \ell)$ times.

Lastly, we analyze the total cost incurred by all inner iterations, which is the cost for increasing the subset connectivity of a graph by one. We may apply Theorem 2.1 directly to analyze the cost of the solution. However, this leads to a bound slightly weaker than what we claimed. To get the desired bound, we apply a stronger version of Nutov's theorem [21]. In particular, the approximation guarantee of Nutov's algorithm depends on the size of a smallest deficient set. To be precise, the size of a smallest deficient set is defined by

$$\min\{|U \cap T| : U \text{ is a deficient set}\}.$$

Lemma 3.12 (Nutov 2009 [21]). *Consider the problem of increasing the rooted subset connectivity of a graph from ℓ to $\ell + 1$. Let $\phi = \min\{|U \cap T| : U \text{ is a deficient set}\}$. That is, each deficient set of the initial graph contains at least ϕ terminals. Then there is an $O(\ell/\phi)$ -approximation algorithm.*

Now, we analyze the size of a smallest deficient set of a graph at the beginning of each inner iteration. Consider the cores at any inner iteration. We call cores at the beginning of the iteration *old cores* and call cores at the end of the iteration *new cores*. We claim that every new core \hat{C}

contains at least two old cores C and D that are disjoint on T . This follows from the following lemma.

Lemma 3.13. *No small deficient set contains two distinct cores C and D such that $C \cap D \cap T \neq \emptyset$.*

Proof. Suppose to the contrary that there is a small deficient set U that contains two distinct cores C and D such that $C \cap D \cap T \neq \emptyset$. Since $C \cap D \cap T \neq \emptyset$, by Lemma 3.4, $C^* \cap D^* \cap T = \emptyset$; that is, $C^* \cap D^*$ has no terminals. Since U contains both C and D , it follows that U^* is contained in both C^* and D^* . Hence, U^* has no terminals. This contradicts the fact that U is a deficient set. \square

Lemma 3.13 implies that no new cores contain two old cores that are intersecting on T . This is because new cores are small deficient sets of the old graph. Moreover, since all small deficient sets that contain only one core have been covered, new cores must contain at least two old cores that are disjoint on T . Thus, the size of a smallest deficient set increases by a factor of 2. This implies the following lemma.

Lemma 3.14. *Consider the j -th inner iteration. At the beginning of the iteration, the size of a smallest deficient set of the current graph is at least 2^{j-1} .*

Proof. As in the above discussion, the size of a smallest deficient set increases by a factor of two in each inner iteration. In more detail, consider the size of a smallest deficient set at the beginning and the end of an inner iteration. We call the graph at the beginning of the iteration an “old graph” and the graph at the end of the iteration a “new graph”. Let U and U' denote smallest deficient sets of the old and the new graph, respectively. By the size argument, we conclude that U and U' are cores of the old and the new graph. At the end of the inner iteration, small deficient sets containing one core are all covered. Thus, U' contains two distinct cores C and C' of the old graph. Moreover, Lemma 3.13 implies that C and C' have no terminals in common. By the minimality of U , we have $|C \cap T| \geq |U \cap T|$ and $|C' \cap T| \geq |U \cap T|$. Thus, $|U' \cap T| \geq |C \cup T| + |C' \cap T| \geq 2|U \cap T|$ as claimed.

Now, we prove the lemma by induction. At the first inner iteration, each deficient set contains at least one terminal. Thus, the statement holds for the base case. Assume that the assertion is true for the $(j-1)$ -th inner iteration; that is, at the beginning of the $(j-1)$ -th inner iteration, any deficient set U has at least 2^{j-2} terminals. By the above claim, this number increases by a factor of two at the end of the iteration. Thus, at the beginning of the j -th iteration, the size of a smallest deficient sets is 2^{j-1} , proving the lemma. \square

By Lemma 3.12 and 3.14, at the j -th inner iteration, the cost incurred by the rooted subset $(\ell+1)$ -connectivity algorithm is $O(\ell/2^{j-1})$. Combining everything together, the approximation guarantee for the problem of increasing the subset connectivity of a graph by one is

$$O\left(\frac{\ell}{2^0} \log \ell + \frac{\ell}{2^1} \log \ell + \dots\right) = O(\ell \log \ell).$$

Thus, by Theorem 3.1, our algorithm achieves an approximation guarantee of $O(k \log^2 k)$, assuming that $|T| \geq 2k$.

3.6 An $O(k \log k)$ -approximation algorithm for $|T| \geq k^2$.

To finish, we show that if the number of terminals is large, then we get a slightly better performance guarantee. Observe that if $|T| \geq k^2$, then, by Lemma 3.10, there is a terminal \hat{r} with a thickness of at most $\frac{q\ell}{|T|} \leq \frac{2\ell^2}{\ell^2} = 2$. Moreover, by Lemma 3.5, each terminal is contained in at most $\frac{2|T|}{|T|-\ell} = O(1)$ halo-sets. Thus, the number of halo-families that are not hit by \hat{r} is $O(1)$. This means that we can hit all the remaining halo-families by choosing $O(1)$ terminals; that is, for each halo-family, we choose one terminal from its core. So, we can skip the micro iterations of the covering-procedure, and the approximation guarantee becomes $O(k \log k)$.

3.7 Analysis for the case $k < |T| < 2k$

Our algorithm in Section 3.5 indeed applies to the case $k < |T| < 2k$ with an approximation guarantee of $O\left(\left(\frac{|T|}{|T|-k}\right)^2 k \log^2 k\right)$. To see this, we leave the bounds in Lemma 3.5 and Lemma 3.10 untouched. Then we have

- Each terminal is contained in at most $O\left(\frac{|T|}{|T|-k}\right)$ halo-families.
- There is a terminal with a thickness of $O\left(\frac{\ell q}{|T|}\right)$, where q is the number of halo-families.

Recall the micro iterations of the covering-procedure. In each micro iteration, we choose $O\left(\frac{|T|}{|T|-\ell}\right)$ terminals, and the number of halo-families (which are not hit) decreases by a factor of $\frac{|T|}{\ell}$. Here the number of micro iterations is not logarithmic because $\frac{|T|}{\ell}$ is not constant when $|T| \approx \ell$. To analyze the upper bound, we write $\frac{\ell}{|T|}$ as $1 - \frac{1}{|T|/(|T|-\ell)}$ and apply an equation:

$$\lim_{x \rightarrow \infty} \left(1 - \frac{1}{x}\right)^x = \frac{1}{e}$$

Thus, we need $O\left(\frac{|T|}{|T|-\ell}\right)$ micro iterations to decrease the number of halo-families (which are not hit) by a factor of e . This means that the covering-procedure terminates in $O\left(\frac{|T|}{|T|-\ell} \log q\right)$ iterations, where q is the number of halo-families. (Note that, in this case, we do not need the preprocessing step because the number of halo-families is at most $|T|^2 = O(k^2)$.) So, the covering-procedure has to call the rooted subset $(\ell + 1)$ -connectivity algorithm for $O\left(\left(\frac{|T|}{|T|-\ell}\right)^2 \log k\right)$ times. Following the analysis in Section 3.5.1, we have an approximation guarantee of $O\left(\left(\frac{|T|}{|T|-k}\right)^2 k \log^2 k\right)$ as claimed.

4 Hardness of the subset k -connectivity problem

In this section, we discuss the hardness of the subset k -connectivity problem. First, we will show in Section 4.1 that the hardest instance of the subset k -connectivity problem might be when $k \approx |T|$; that is, we prove Proposition 2.3. Then we will present in Section 4.2 an approximation preserving reduction from the rooted subset k -connectivity problem to the subset k -connectivity problem; that is, we prove Theorem 2.4.

4.1 The hardest instance

We will show that an $\alpha(k)$ -approximation algorithm for the case $|T| = k$ implies an $(\alpha(k) + f(k))$ -approximation algorithm for all instances with $|T| > k$, where $f(k)$ is the best known approximation guarantee for the rooted subset k -connectivity problem. In particular, instances with $|T| \approx k$ might be the hardest cases of the subset k -connectivity problem.

Suppose there is an $\alpha(k)$ -approximation algorithm \mathcal{A} for the subset k -connectivity problem for the case $|T| = k$. We apply \mathcal{A} to solve an instance of the subset k -connectivity problem with $|T| > k$ as follows. Let $G = (V, E)$ be a given graph and $T \subseteq V$ be a set of terminals, where $|T| > k$. First, we take any subset R of k terminals from T . Then we apply the algorithm \mathcal{A} to this instance with R as the set of terminals; this results in a graph $G_R = (V, E_R)$. Clearly, R is k -connected in G_R . Now, we make the remaining terminals connected to R by applying the rooted subset k -connectivity algorithm. To be precise, we construct a padded graph by adding a new vertex \hat{r} and new edges of zero cost from \hat{r} to each terminal of R . Then we apply the rooted subset k -connectivity algorithm to the padded graph with the set of terminals T and the root \hat{r} . Denote a solution subgraph (of the padded graph) by $G_{pad} = (V \cup \{\hat{r}\}, E_{root} \cup \{(\hat{r}, t) : t \in R\})$, where $E_{root} \subseteq E$. The algorithm outputs the union of the two subgraphs, namely $\hat{G} = (V, E_R \cup E_{root})$.

We claim that the set of all terminals T is k -connected in \hat{G} . Suppose not. Then there is a set of vertices $X \subseteq V$ of size $k - 1$ that separates some terminals $s, t \in T - X$; that is, s and t are not connected in $\hat{G} - X$. Consider the padded subgraph G_{pad} . By the construction, since G_{pad} is k -connected from \hat{r} to T , both s and t have paths to \hat{r} in $G_{pad} - X$. Moreover, each of these two paths must visit some terminals s' and t' in R , respectively. If $s' = t'$, then s and t are connected by the union of these paths. So, we have a contradiction. If $s' \neq t'$, then we can join these two paths by an s', t' -path in $G_R - X$. Such s', t' -path exists because R is k -connected in G_R , meaning that X cannot separate a pair of terminals in R . Thus, s and t are connected, and we again have a contradiction.

Now, consider the cost. The approximation factor incurred by the algorithm \mathcal{A} is $\alpha(k)$, and the approximation factor incurred by the rooted subset k -connectivity algorithm is $f(k)$. Thus, the above algorithm gives an approximation guarantee of $(\alpha(k) + f(k))$ as claimed.

4.2 A reduction from the rooted subset k -connectivity problem

As we showed in the previous section, an approximation algorithm for the rooted subset k -connectivity problem implies an approximation algorithm for the subset k -connectivity problem. Hence, it is more likely that the rooted problem is easier than the subset problem. Here we show a solid evidence of this statement; that is, we will give an approximation preserving reduction from the rooted subset k -connectivity problem to the subset k -connectivity problem.

The key idea of the reduction is that a solution \hat{G} to the rooted subset k -connectivity problem is indeed almost subset k -connected. In particular, if the root vertex r is not allowed to be removed, then there is no set of vertices of size less than k that can separate a pair of terminals. So, we want to prevent the root vertex r from being in a separator. To do this, we replace r by a clique K_d of size at least $k + 1$. Thus, removing any set of less than k vertices cannot remove all vertices corresponding to r .

Now, we shall realize the above idea. First, take any instance Π^{root} of the rooted subset k -connectivity problem consisting of a graph $G = (V, E)$, a set of terminals $T \subseteq V$ and a root vertex $r \in V - T$. Let d be the degree of r in G . Clearly, if the instance Π^{root} is feasible, then

$d \geq k$. We construct an instance Π^{subset} of the subset k -connectivity problem consisting of a graph $G' = (V', E')$ and a set of terminals T' as follows. Let $\{v_1, v_2, \dots, v_d\}$ be a set of neighbors of r in G . We remove from G the vertex r and replace it with a clique K_{d+1} on a set of vertices $\{r', v'_1, v'_2, \dots, v'_d\}$. All edges of K_{d+1} have zero costs. The vertex r' corresponds to the root vertex r of G , and each vertex v'_i corresponds to each neighbor v_i of r in G . Then we connect K_{d+1} to G by adding to G' an edge (v'_i, v_i) for each edge (r, v_i) in G and setting the cost of (v'_i, v_i) to be the same as the cost of (r, v_i) . Thus, each edge (v'_i, v_i) in G' corresponds to an edge (r, v_i) in G . The set of terminals of this new instance is $T' = T \cup \{r'\}$, and the connectivity requirements is k , the same for both instances. This completes the construction.

In sum, we have

$$\begin{aligned} G' &= G - \{r\} + K_{d+1} + \{(v'_1, v_1), (v'_2, v_2), \dots, (v'_d, v_d)\} \\ T' &= T \cup \{r'\} \\ r' &\leftrightarrow r \\ (v'_i, v_i) &\leftrightarrow (r, v_i) \text{ for all } i = 1, 2, \dots, d \end{aligned}$$

Completeness: First, we show that any feasible solution H of Π^{root} maps to a feasible solution H' of Π^{subset} with the same cost. The mapping is as follows. Given a graph H , we construct a solution H' to Π^{subset} by taking all edges of K_{d+1} and all edges of G' corresponding to edges of H . Clearly, the cost of H' and H are the same. It remains to show that T' is k -connected in H' .

The connectivity between the vertex r' and each terminal $t \in T$ is clearly satisfied. This is because any collection of openly disjoint r, t -paths in H maps to a collection of openly disjoint r', t -paths in H' . In particular, any path $P = (r, v_i, \dots, t)$ in H maps to a path $P' = (r, v'_i, v_i, \dots, t)$ in H' , and it is easy to see that the mapping preserves vertex-disjointness. By the same argument, we can deduce that every vertex $v'_j \in K_{d+1}$ is k -connected to t in H' . This is because the path P also maps to a path $P'' = (v'_j, v'_i, v_i, \dots, t)$ or $P'' = (v'_j, v_j, \dots, t)$ in H' .

Now, consider the connectivity between a pair of vertices $t, t' \in T$. Assume a contradiction that t and t' are not k -connected. Then there is a subset of vertices X of G' with $|X| \leq k - 1$ such that t and t' are not connected in $H' - X$. Since $|X| \leq k - 1 \leq d$, there is a vertex s in $K_{d+1} - X$. (The vertex s is either the vertex r' or some vertex v_i in K_{d+1} .) As we have shown, s is k -connected to t and t' in H' . Thus, by Menger's theorem, $H' - X$ contains both an s, t -path and an s, t' -path. So, t and t' are connected in $H' - X$, a contradiction. Therefore, T' is k -connected in H' , implying that H' is feasible to the subset k -connectivity problem.

Soundness: Now, we show the converse; that is, any feasible solution H' of Π^{subset} maps to a feasible solution H of Π^{root} with the same cost. This direction is easy. We construct H by taking all edges of H' that correspond to edges of G . Clearly, the cost of H and H' are the same. By feasibility, H' has, for each terminal $t \in T$, a collection of k openly disjoint r', t -paths, namely P'_1, P'_2, \dots, P'_k , and each path P'_j is of the form $P'_j = (r', v'_i, v_i, \dots, t)$. The path P'_j maps to a path $P_j = (r, v_i, \dots, t)$ in H . So, we have a collection of paths P_1, P_2, \dots, P_k in H that are openly disjoint. Therefore, H is feasible to the rooted subset k -connectivity problem, finishing the proof.

5 Conclusions and Discussions

We studied the structure of the subset k -connectivity problem and used this knowledge to design an approximation algorithm for the subset k -connectivity problem. When the number of terminals is moderately large, at least $2k$, our algorithm gives a very good approximation guarantee of $O(k \log^2 k)$. When the number of terminals is tiny, at most \sqrt{k} , then the trivial algorithm also gives a very good approximation guarantee of $O(k)$. However, when the number of terminals is between \sqrt{k} and $2k$, the approximation guarantee can be as large as $\Theta(k^2)$. Interestingly, as we have shown, it does seem that the hardest instances of the subset k -connectivity problem are when the number of terminals is close to k .

Acknowledgments. We thank Joseph Cheriyan for useful discussions over a year. Also, we thank Adrian Vetta, Parinya Chalermsook, Danupon Nanongkai, Jittat Fakcharoenphol and anonymous referees for useful comments on the preliminary draft.

References

- [1] Vincenzo Auletta, Yefim Dinitz, Zeev Nutov, and Domenico Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *J. Algorithms*, 32(1):21–30, 1999.
- [2] Jaroslav Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. An improved LP-based approximation for steiner tree. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*, pages 583–592, 2010.
- [3] Tanmoy Chakraborty, Julia Chuzhoy, and Sanjeev Khanna. Network design for vertex connectivity. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 167–176, 2008.
- [4] Joseph Cheriyan, Bundit Laekhanukit, Guylain Naves, and Adrian Vetta. Approximating rooted steiner networks. In *SODA*, pages 1499–1511, 2012.
- [5] Joseph Cheriyan, Santosh Vempala, and Adrian Vetta. An approximation algorithm for the minimum-cost k -vertex connected subgraph. *SIAM J. Comput.*, 32(4):1050–1055, 2003.
- [6] Joseph Cheriyan and Adrian Vetta. Approximation algorithms for network design with metric costs. *SIAM J. Discrete Math.*, 21(3):612–636, 2007.
- [7] Julia Chuzhoy and Sanjeev Khanna. Algorithms for single-source vertex connectivity. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*, pages 105–114, 2008.
- [8] Julia Chuzhoy and Sanjeev Khanna. An $O(k^3 \log n)$ -approximation algorithm for vertex-connectivity survivable network design. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09)*, pages 437–441, 2009.
- [9] Jittat Fakcharoenphol and Bundit Laekhanukit. An $O(\log^2 k)$ -approximation algorithm for the k -vertex connected spanning subgraph problem. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 153–158, 2008.

- [10] Lisa Fleischer, Kamal Jain, and David P. Williamson. Iterative rounding 2-approximation algorithms for minimum-cost vertex connectivity problems. *J. Comput. Syst. Sci.*, 72(5):838–867, 2006.
- [11] Michel X. Goemans, Andrew V. Goldberg, Serge A. Plotkin, David B. Shmoys, Éva Tardos, and David P. Williamson. Improved approximation algorithms for network design problems. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 223–232, 1994.
- [12] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001.
- [13] Samir Khuller and Balaji Raghavachari. Improved approximation algorithms for uniform connectivity problems. *J. Algorithms*, 21(2):434–450, 1996.
- [14] Guy Kortsarz, Robert Krauthgamer, and James R. Lee. Hardness of approximation for vertex-connectivity network design problems. *SIAM J. Comput.*, 33(3):704–720, 2004.
- [15] Guy Kortsarz and Zeev Nutov. Approximating k -node connected subgraphs via critical graphs. *SIAM J. Comput.*, 35(1):247–257, 2005.
- [16] Bundit Laekhanukit. Approximation algorithms for (S, T) -connectivity problems. Master’s thesis, University of Waterloo, Canada, 2010.
- [17] Zeev Nutov. An almost $O(\log k)$ -approximation for k -connected subgraphs. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 912–921, 2009.
- [18] Zeev Nutov. A note on rooted survivable networks. *Inf. Process. Lett.*, 109(19):1114–1119, 2009.
- [19] Zeev Nutov. Approximating steiner networks with node-weights. *SIAM J. Comput.*, 39(7):3001–3022, 2010.
- [20] Zeev Nutov. Approximating subset k -connectivity problems. In *Proceedings of the 9th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 9–20, 2011.
- [21] Zeev Nutov. Approximating minimum-cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms*, 9(1):1, 2012.
- [22] R. Ravi and David P. Williamson. An approximation algorithm for minimum-cost vertex-connectivity problems. *Algorithmica*, 18(1):21–43, 1997.